



IBM Software Group | Tivoli Software

# Managing Linux for Control and Performance

***Laura Knapp***

***ljknapp@us.ibm.com***



**ON DEMAND BUSINESS™**

© 2005 IBM Corporation

# Agenda

- **Introduction and background**
- **Performance Methodologies**
- **6 Common Problems**
  - **Disk Utilization?**
  - **What's going on in the IP stack?**
  - **What's going on in TCP and UDP?**
  - **What are response times?**
  - **What is system availability?**
  - **User Activity**



# Increasing Importance of Performance

## Performance Management

The practice of managing network service response time, consistency and quality for individual services and services overall

## Performance Related Risks

Network degradation and failure

Application timeouts and failure

Application degradation



## Loss of Customers



# The Performance Problem

## Over-provisioning

- Lots of provisions (rare)
- More resources than can be consumed
  - Food on a cruise
  - Congressional parking spaces
  - AOL CD-ROMs

## Over-subscribing - lots of subscribers

- Lots of subscribers (common)
- Many users consume all the resources
  - Batteries, chain saws, interstate lanes during a hurricane
  - Phone calls on Mothers' Day
  - Many to few: whenever there's a bottleneck or funnel
  - Fast to slow: things will back up

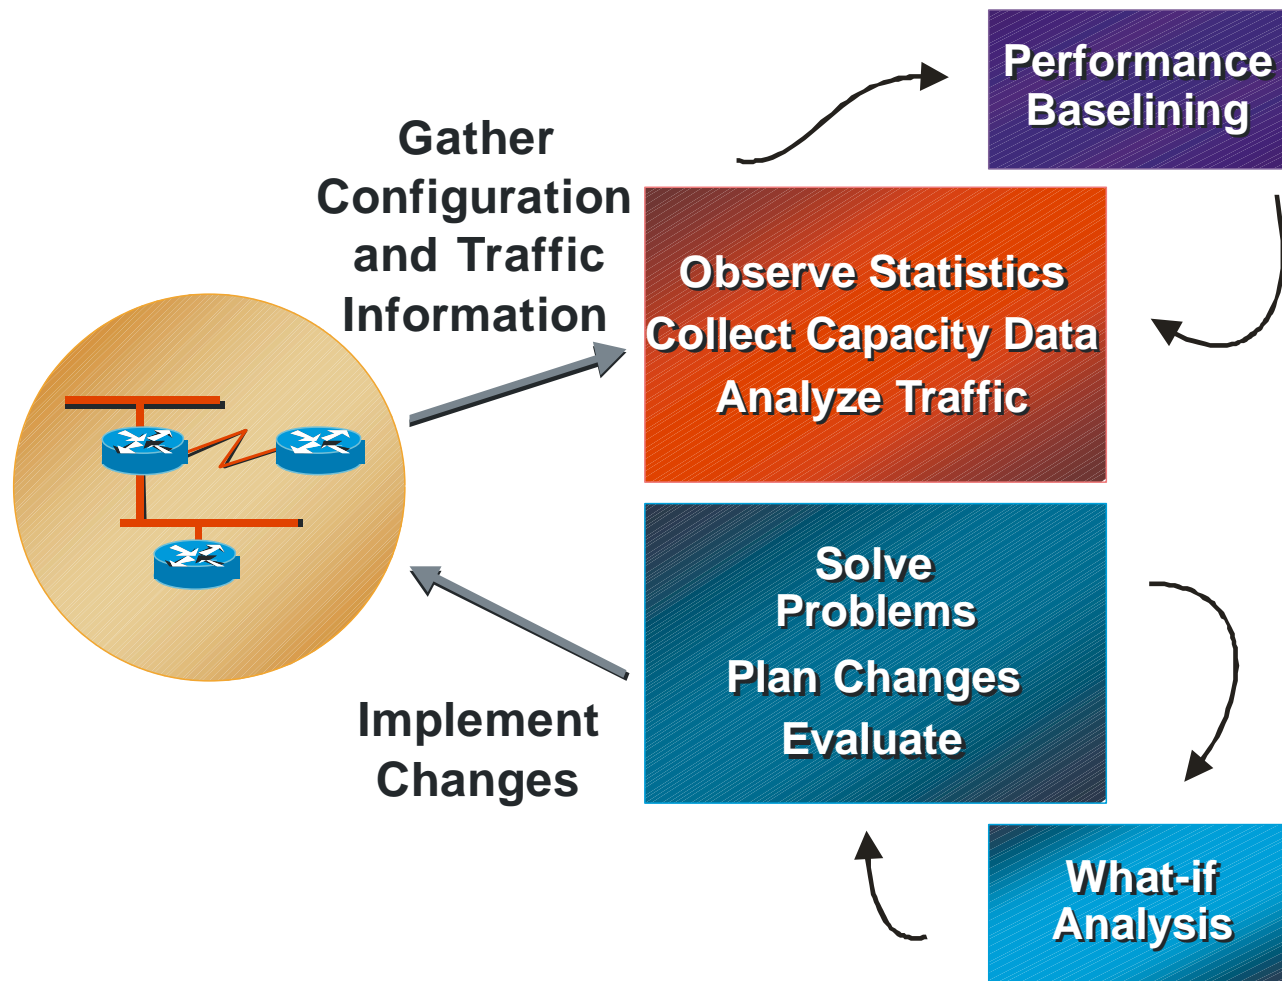


# Agenda

- Introduction and background
- **Performance Methodologies**
- 6 Common Problems
  - Disk Utilization?
  - What's going on in the IP stack?
  - What's going on in TCP and UDP?
  - What are response times?
  - What is system availability?
  - User Activity
  -



# Effective Performance Management



# IP Resource Bottlenecks

**CPU**

**Memory**

**Buffering, queuing, and latency**

**Interface and pipe sizes**

**Network Capacity**

**Speed and distance**


**Application characteristics**

**Results in**

**Network capacity problems**

**Utilization overload**

**Application failure**





# Performance Plan

## Develop information collection plan

Define parameters to be monitored/measured and the threshold

Acquire proper authority to change threshold

Determine frequency of monitoring and reporting

Determine frequency of alerting mechanism

Define parameters that trigger alert mechanism

Define performance areas of interest

Report and interpret results

Determine tools for collecting information





## Agenda

- Introduction and background
- Performance Methodologies
- **6 Common Problems**
  - Disk Utilization?
  - What's going on in the IP stack?
  - What's going on in TCP and UDP?
  - What are response times?
  - What is system availability?
  - User Activity
  -




## Disk Utilization

**First line of information slowdown is if the storage device cannot move the data**

**Determining what is currently being used on the system will assist in determining how much you can grow the system**

**An application behaving poorly may be due to improper design, improper setting of system resources to use, or application configuration**

**Sluggishness of a system may be due to not enough CPU, I/O overloads, or queue latencies**



## Key Items Disk Utilization

**What Disks are available?**

**What is the data usage of each?**

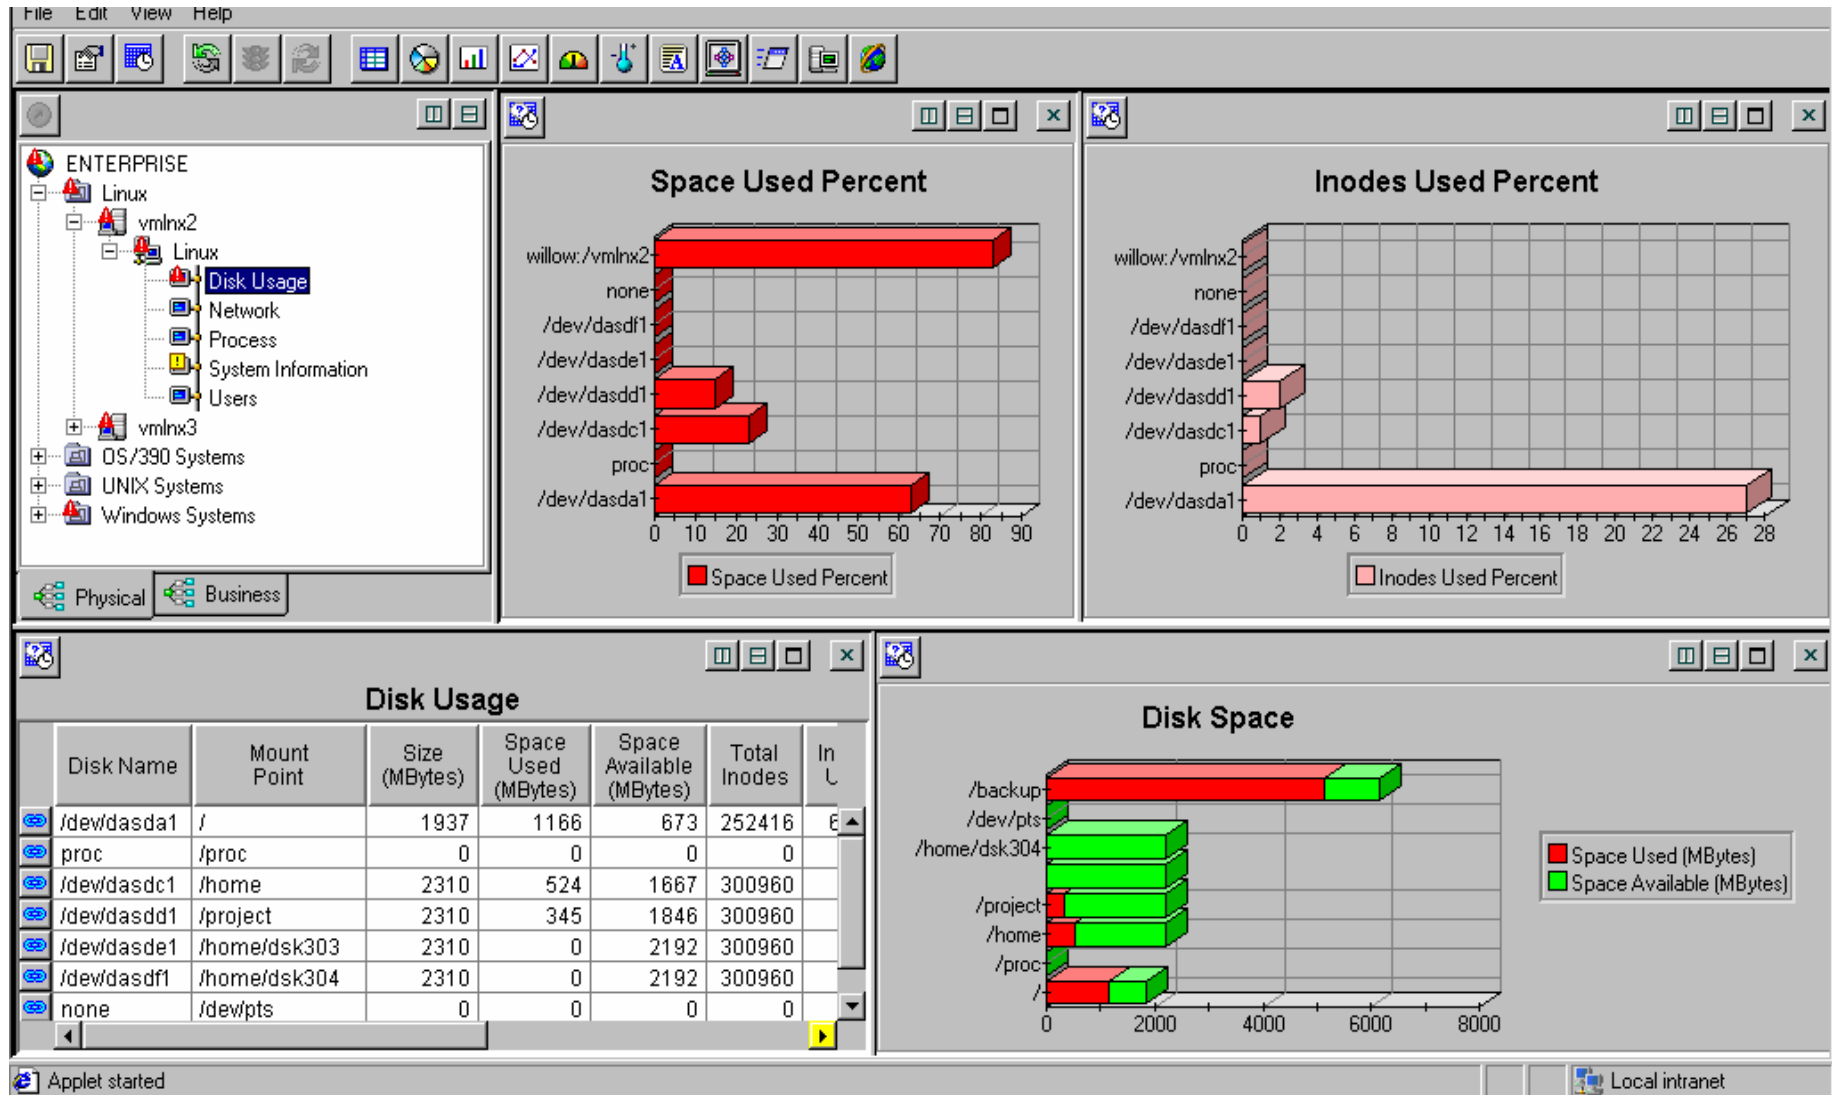
**What is the inodes usage?**

**What type of device is available?**

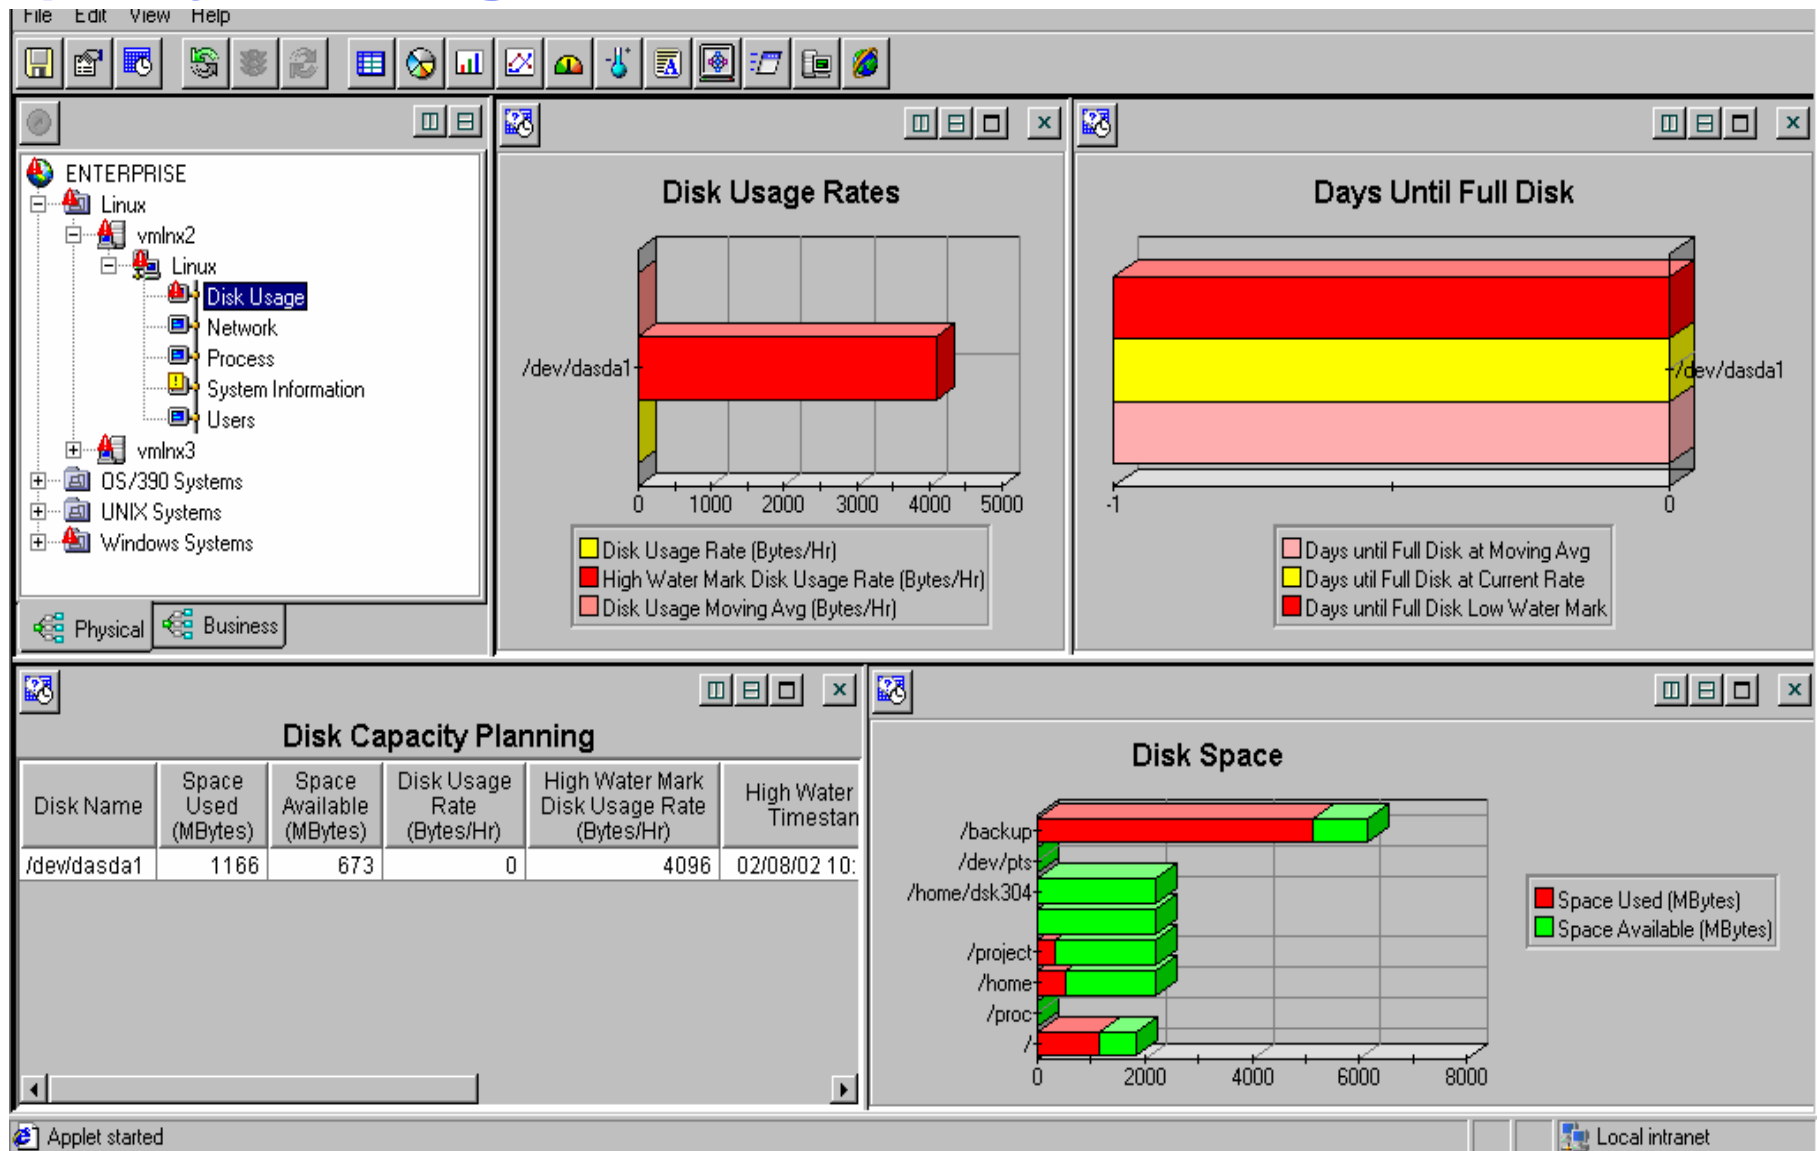
**What is the current status?**



# How's the Disk Space?

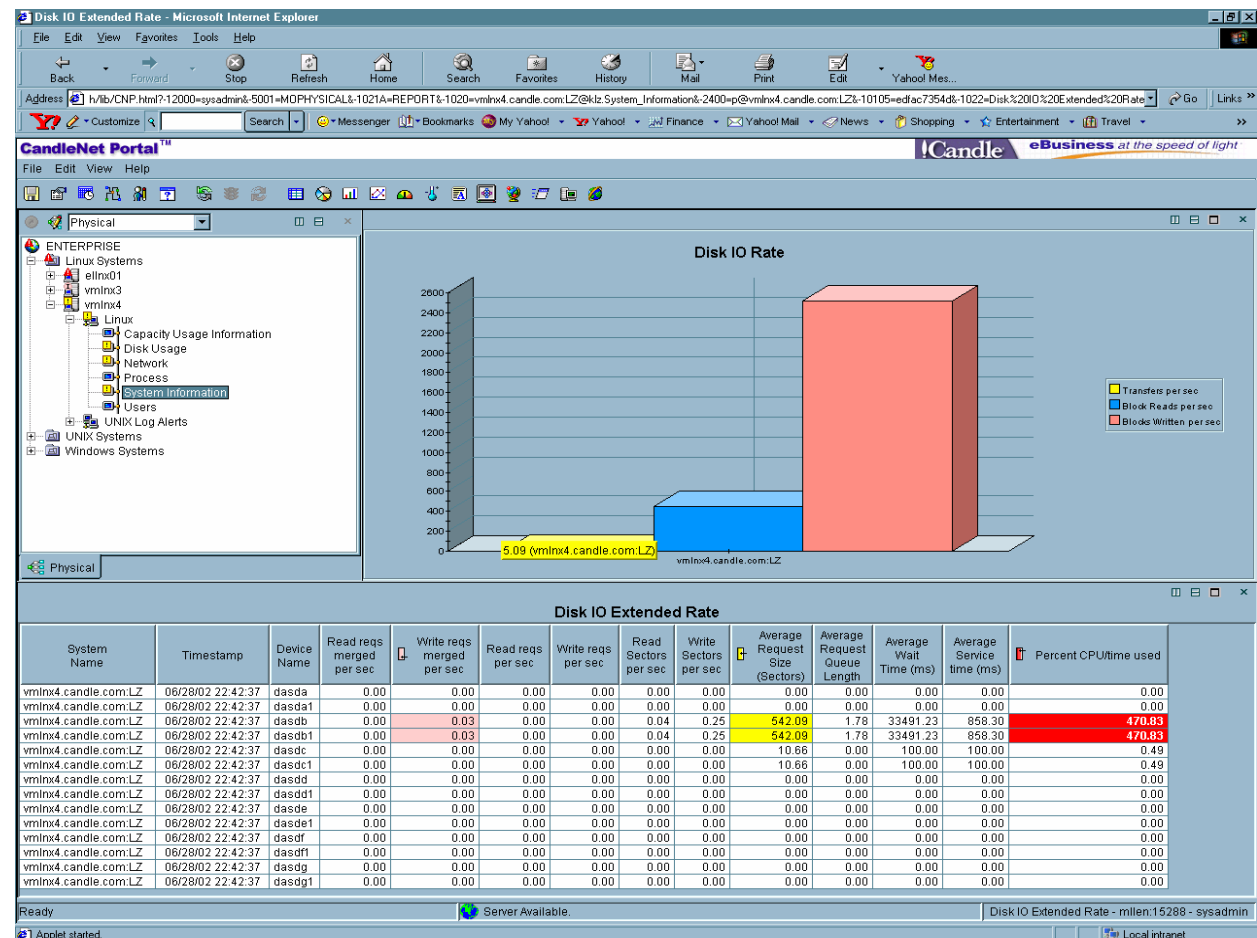


# Capacity Planning for Disk Utilization



# Extended I/O Information for 2.4 Kernels

- More extended I/O information available directly in our agent.
- Can be accessed via `iostat -x` command.
- Other tools Sysadmins use are now provided in our agent.



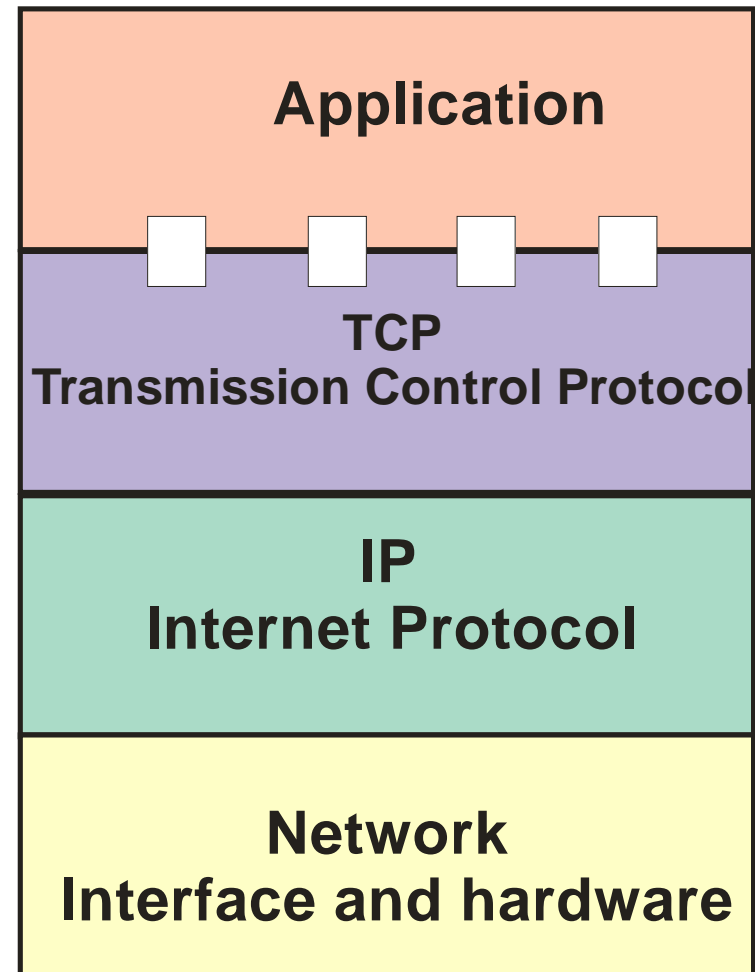
## The IP Stack

**The IP stack is responsible for moving your data through the network. The datagram is the basic unit of transmission**

**Delivery is its role not validation**

**Determines if a datagram needs to be sent to a router in order to reach its destination**

**Prepares the frame for the underlying network**



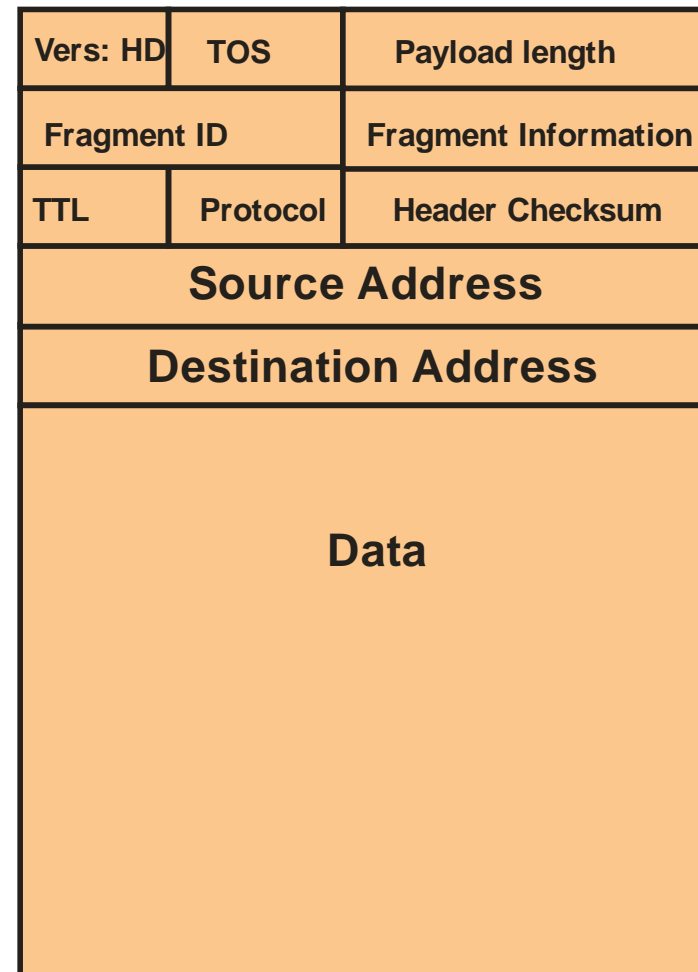


## Understanding the IP Stack

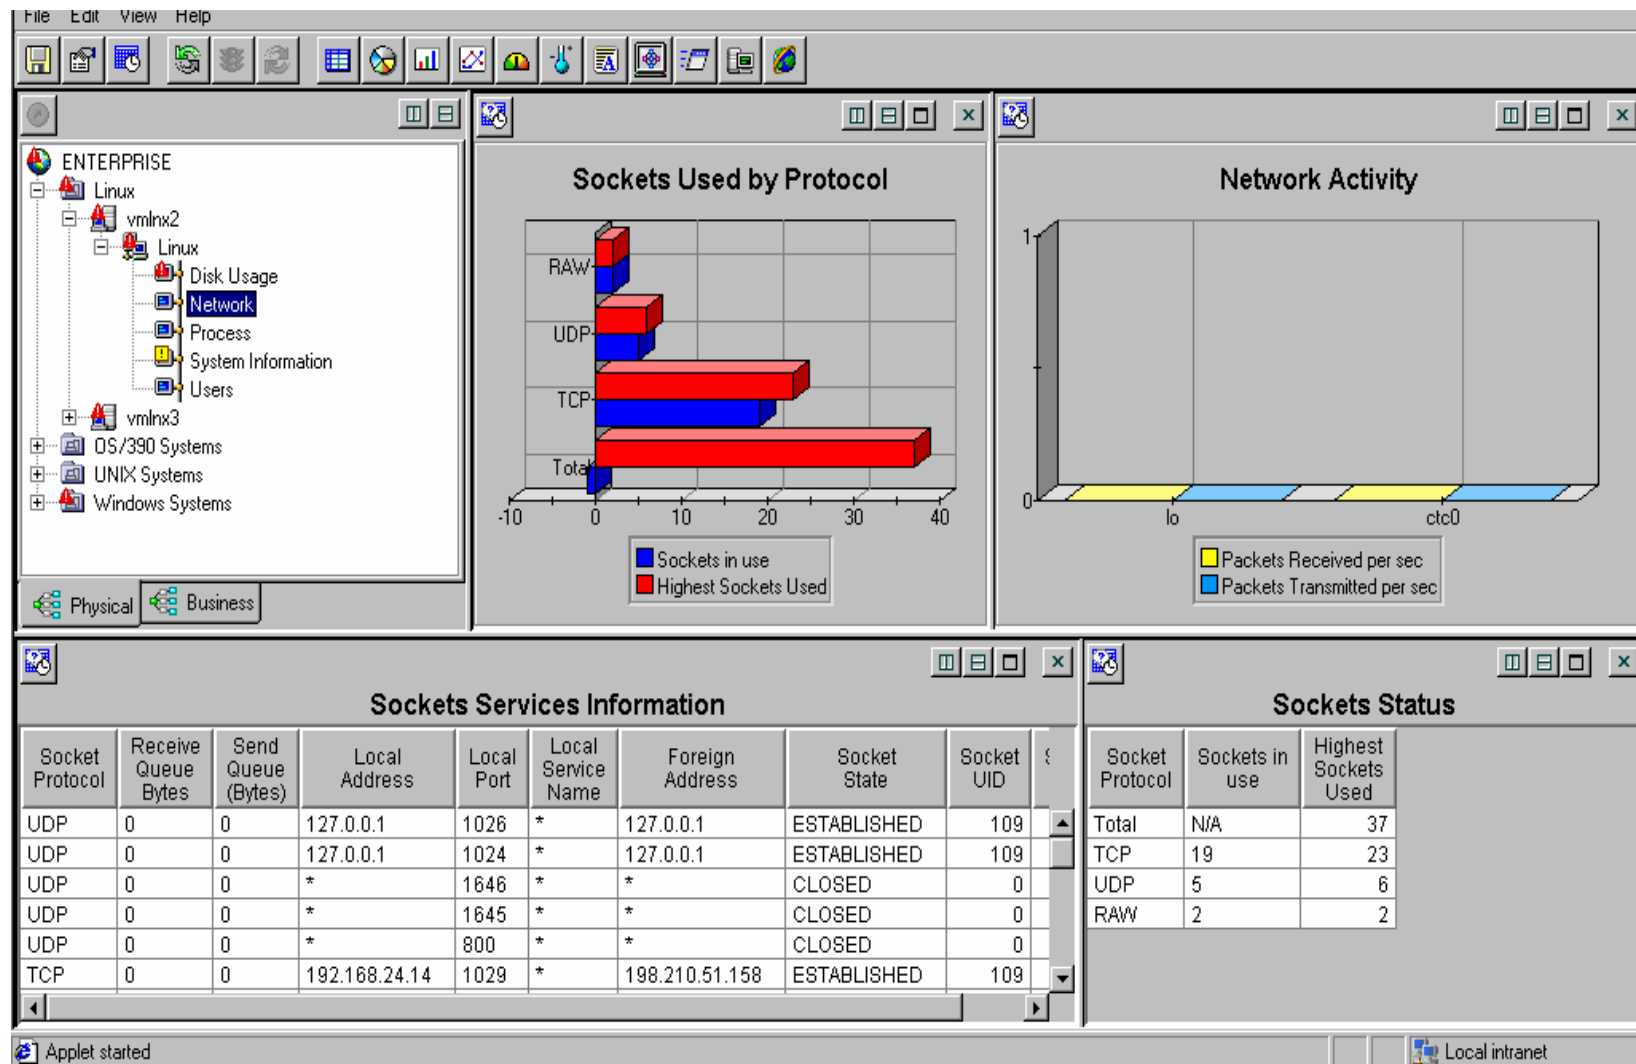
- **Datagram**

**One packet, or unit, of information that includes relevant delivery information, such as the destination address, that is sent through a packet-switching network.**


- **Receives**
- **Forwards**
- **Discards**
- **Retransmission**



# Network Resource Utilization



## IP Stack Data



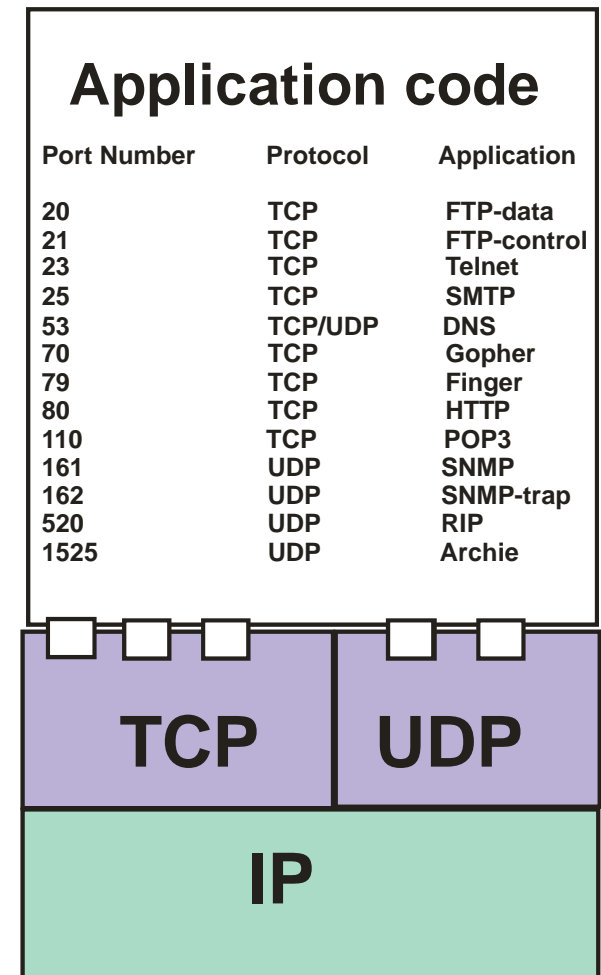
## The UDP Stack

**Application dependent if this layer 4 stack is used**

**Faster because no connections are set up between the two communicating systems**


**Assumes others are making sure that frames arrive in the correct sequence**

**Used by many streaming applications**




## UDP Stack Items of Note

- **Sent**
- **Received**
- **Discarded**
- **Delivered**
- **Not Delivered**
- **Fragmenting**

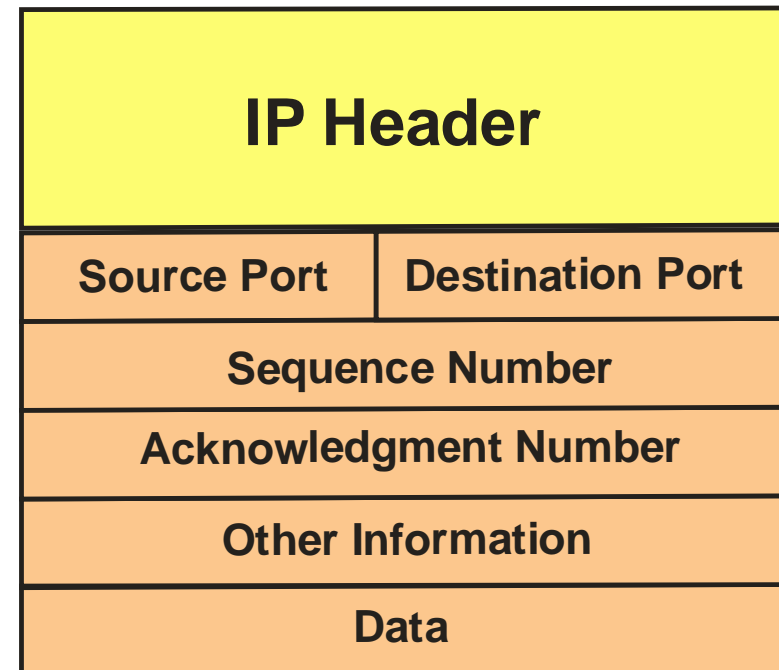


# UDP Stack Information



## TCP Stack Items

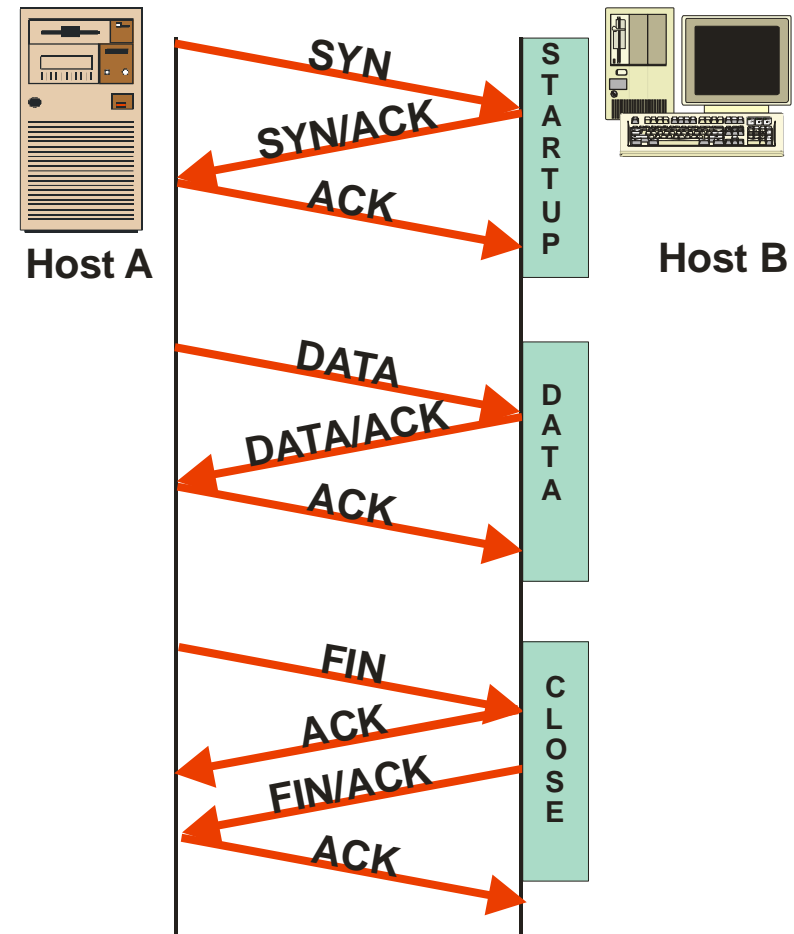
- Most commonly used layer 4 protocol
- Sets up a connection between communicating end points
- Makes sure that all frames are received
- Timers, acknowledgements, are all components
- Sequencing is a component






## TCP Stack Elements

- **Connections**
- **Accepted**
- **Active**
- **Dropped**
- **Segments retransmitted**



# TCP Stack Data



## Response Time

**Web users expect 2 to 5 second response time**

**SNA users expect sub-second response time**

**No one is ever happy with what they get**

**External customers may go elsewhere**

**Where is the problem?**

**Network?**

**Router have long ques?**

**Is the Lan to slow?**

**Is the route long?**

**Operating system?**

**Too long to queue for transmit?**

**Application?**

**Protocol?**

**Window size improperly set?**

**MTU size improperly set?**



## Core Elements of Response Time

**What are overall response times in my network?**

**What are response times for different size frames?**


**Can I look at a specific address and determine its response time?**

**Are both real time and historical views available?**


**Are both graphical and tabular views available?**

**Can I set thresholds?**

**Can I send alerts?**



# Response Time




## Availability

**Resources, applications, network components that are not available impact many aspects of your system**

**IP is especially prone to this due to the 'non-configurable' operations**

**Critical resources can come and go with no 'network-wide' configuration, but this may impact other systems**

**Five steps may occur in a process before you realize that the six step requires a resource that is no longer available**



## Key Items in Availability

**Can you get a view of key application availability?**

**Can you define critical resources?**

**Have alerts been sent?**

**Is the system not available because the system is down or because a resource like a router is having problems or an application is not available?**

**Can I tell if the route is not the normal route taken?**

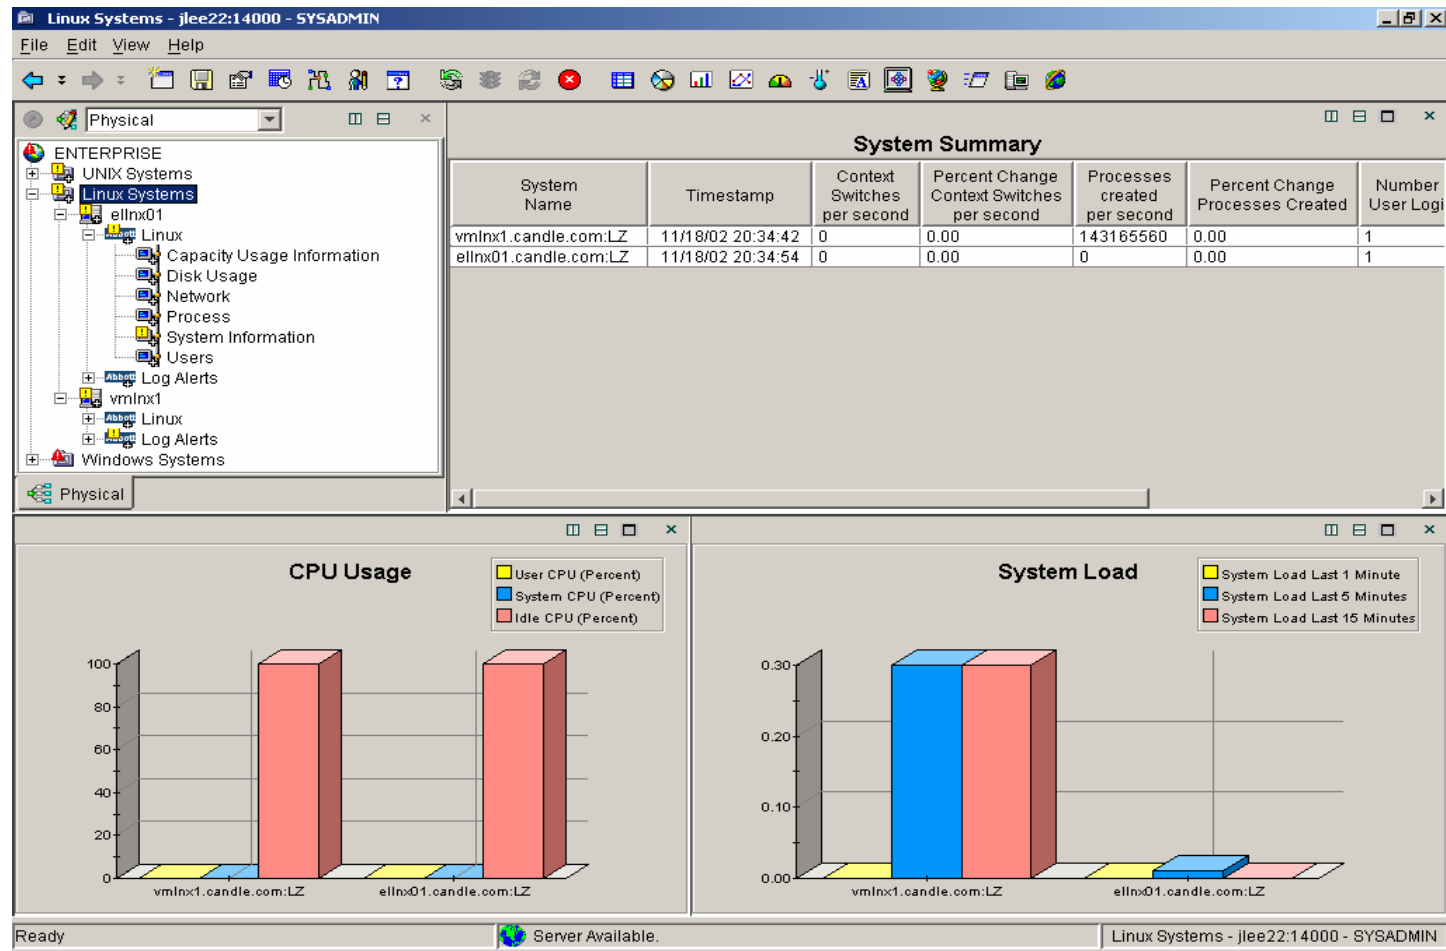
**Has the situation cleared itself up?**

**How can I get more details from an offending intermediary system?**

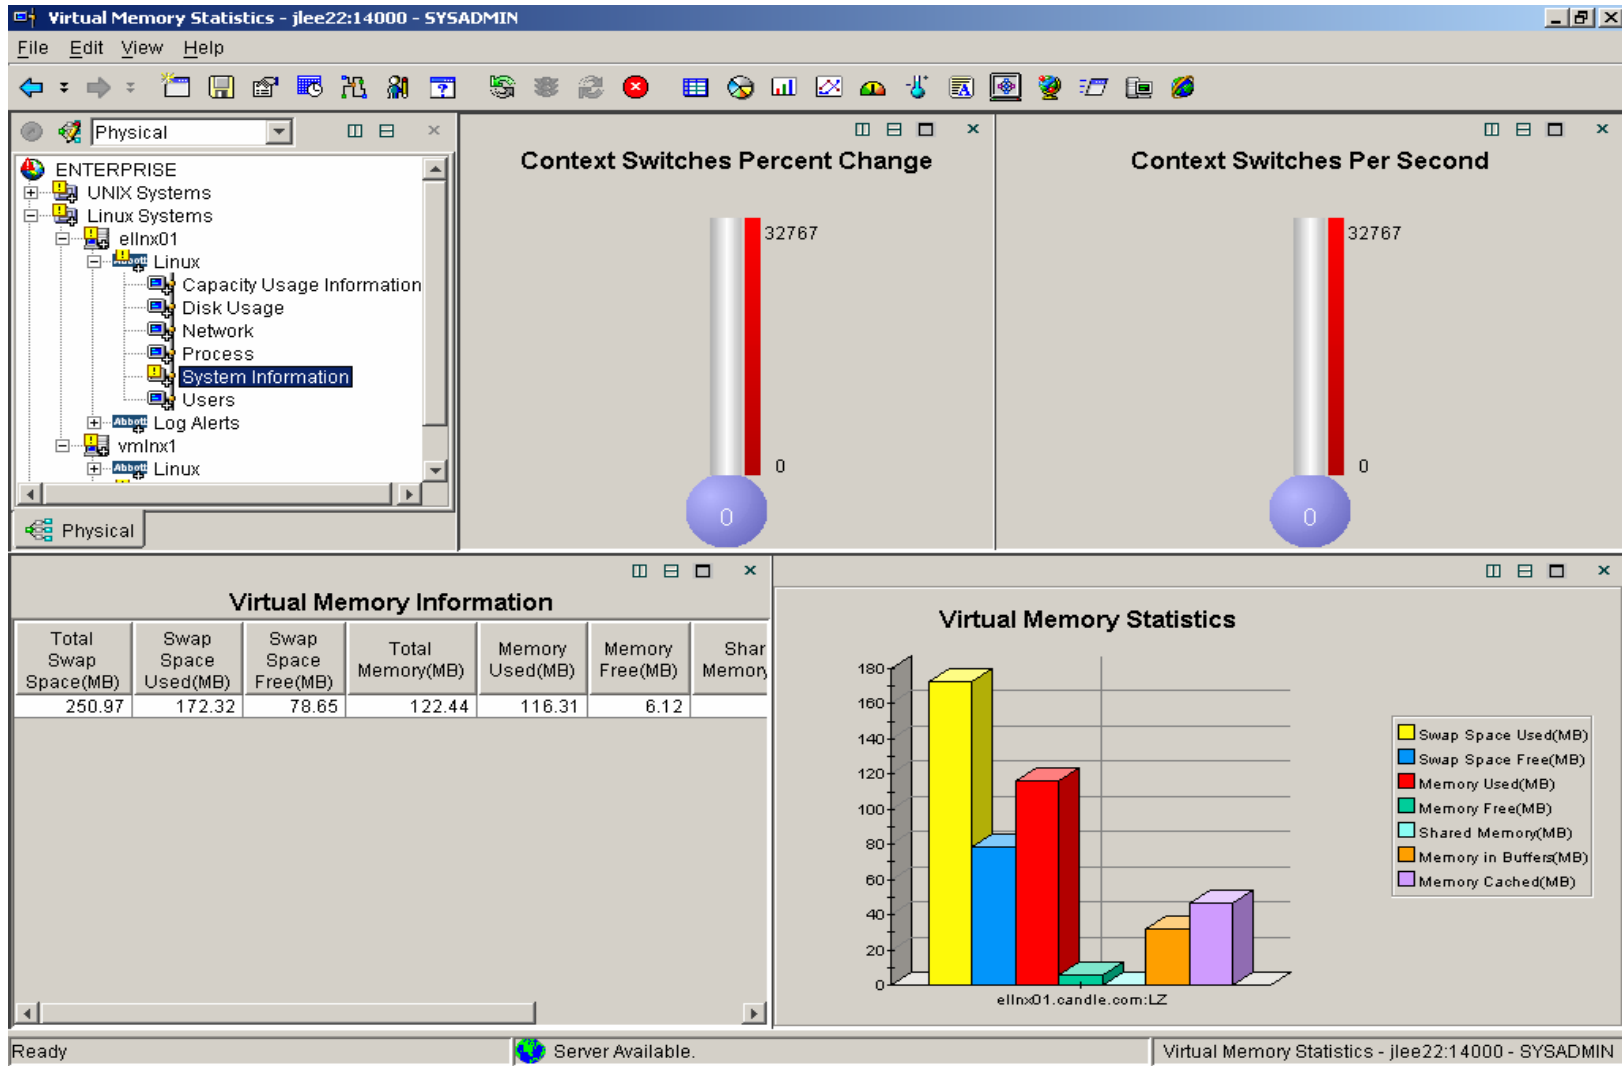




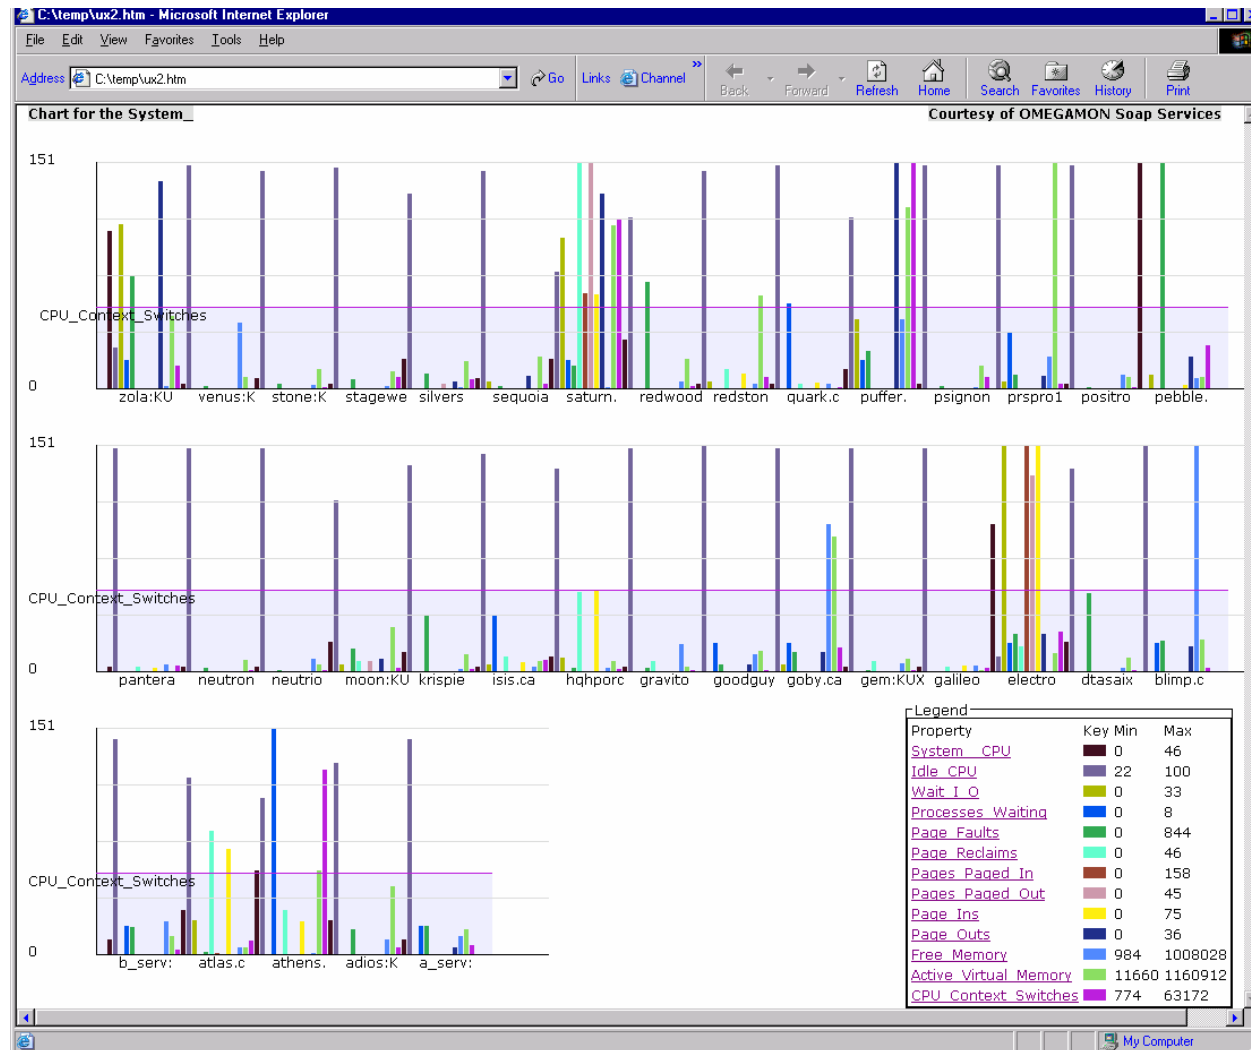
# System Metrics



# Understanding Virtual Memory Usage



# Cross Server Reporting



## Resource Utilization

**Application usage by end users is very unpredictable in IP. What was valid last week may not be valid today**

**An application installed on a system and active not being utilized by end users is taking system resources that could be used by other applications**


**Sometimes it is appropriate to block users after a given number have logged onto an application in order to conserve existing resources**

**Knowing who is using what on a given system can help determine long term capacity planning needs for the system**

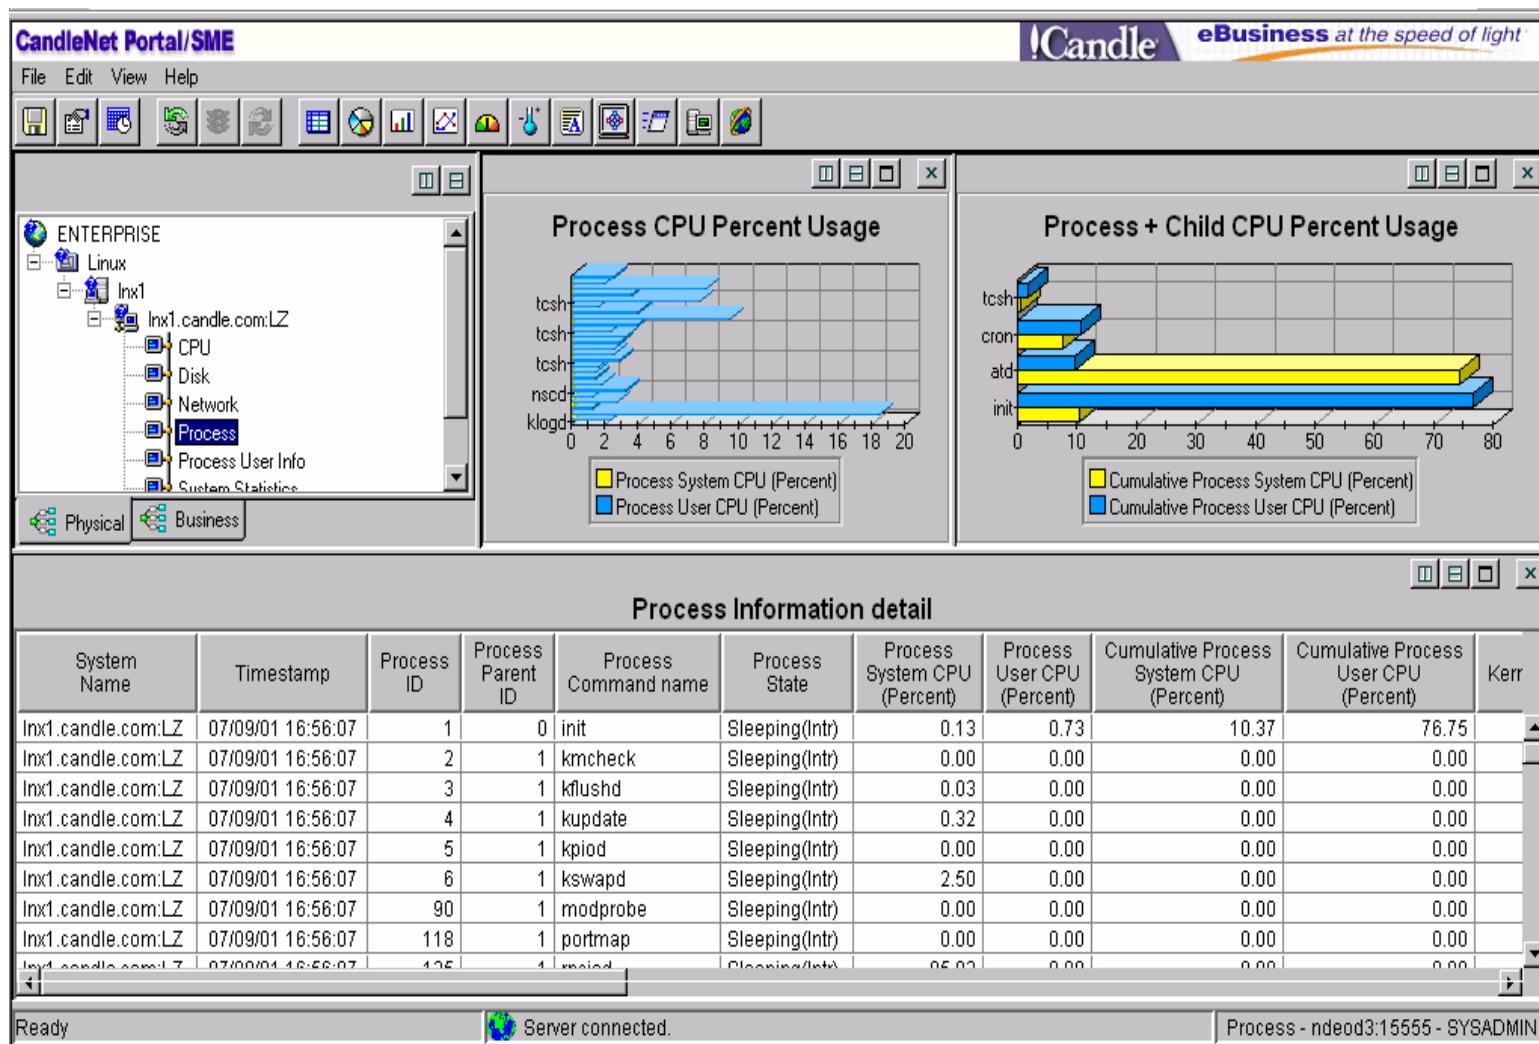


## Resource Utilization Key Items

- For a given system, can you determine the applications being used?
- Can you tell for each application the session or user counts?
- Can you tell for each application the number of bytes transferred?
- Can you get more details on a historical basis?  
Determine patterns
- like a 3% monthly growth in usage?
- Can you alert on miss-use of an application?
- Can you shut a user out of an application in real time?



# Resource Utilization



## Need to See More Detailed Data?

Process Information detail										
System Name	Timestamp	Process ID	Process Parent ID	Process Command name	Process State	Process System CPU (Percent)	Process User CPU (Percent)	Cumulative Process System CPU (Percent)	Cumulative Process User CPU (Percent)	Ken
lnx1.candle.com:LZ	07/09/01 16:56:07	1	0	init	Sleeping(Intr)	0.13	0.73	10.37	76.75	
lnx1.candle.com:LZ	07/09/01 16:56:07	2	1	kmcheck	Sleeping(Intr)	0.00	0.00	0.00	0.00	
lnx1.candle.com:LZ	07/09/01 16:56:07	3	1	kflushd	Sleeping(Intr)	0.03	0.00	0.00	0.00	
lnx1.candle.com:LZ	07/09/01 16:56:07	4	1	kupdate	Sleeping(Intr)	0.32	0.00	0.00	0.00	
lnx1.candle.com:LZ	07/09/01 16:56:07	5	1	kpiod	Sleeping(Intr)	0.00	0.00	0.00	0.00	
lnx1.candle.com:LZ	07/09/01 16:56:07	6	1	kswapd	Sleeping(Intr)	2.50	0.00	0.00	0.00	
lnx1.candle.com:LZ	07/09/01 16:56:07	90	1	modprobe	Sleeping(Intr)	0.00	0.00	0.00	0.00	
lnx1.candle.com:LZ	07/09/01 16:56:07	118	1	portmap	Sleeping(Intr)	0.00	0.00	0.00	0.00	
lnx1.candle.com:LZ	07/09/01 16:56:07	125	1	rpciod	Sleeping(Intr)	95.02	0.00	0.00	0.00	
lnx1.candle.com:LZ	07/09/01 16:56:07	136	1	syslogd	Sleeping(Intr)	0.09	0.73	0.00	0.00	
lnx1.candle.com:LZ	07/09/01 16:56:07	140	1	klogd	Sleeping(Intr)	0.00	2.38	0.00	0.00	
lnx1.candle.com:LZ	07/09/01 16:56:07	187	1	httpd	Sleeping(Intr)	0.35	18.49	0.00	0.00	
lnx1.candle.com:LZ	07/09/01 16:56:07	196	1	atd	Sleeping(Intr)	0.01	0.73	74.29	9.58	
lnx1.candle.com:LZ	07/09/01 16:56:07	203	1	inetd	Sleeping(Intr)	0.01	0.73	0.28	0.47	
lnx1.candle.com:LZ	07/09/01 16:56:07	210	1	lpd	Sleeping(Intr)	0.00	0.00	0.00	0.00	
lnx1.candle.com:LZ	07/09/01 16:56:07	212	187	httpd	Sleeping(Intr)	0.00	0.00	0.00	0.00	
lnx1.candle.com:LZ	07/09/01 16:56:07	245	1	sendmail	Sleeping(Intr)	0.06	1.28	0.00	0.00	
lnx1.candle.com:LZ	07/09/01 16:56:07	249	1	cron	Sleeping(Intr)	0.20	2.74	7.74	10.78	
lnx1.candle.com:LZ	07/09/01 16:56:07	261	1	nscd	Sleeping(Intr)	0.14	1.28	0.00	0.00	
lnx1.candle.com:LZ	07/09/01 16:56:07	262	261	nscd	Sleeping(Intr)	0.32	3.29	0.00	0.00	
lnx1.candle.com:LZ	07/09/01 16:56:07	263	262	nscd	Sleeping(Intr)	0.08	1.83	0.00	0.00	
lnx1.candle.com:LZ	07/09/01 16:56:07	264	262	nscd	Sleeping(Intr)	0.05	1.46	0.00	0.00	

Ready

Server connected.

Process - ndeod3:15555 - SYSADMIN




## User Activity

**You are dealing with a system, not just a standalone computer. Other tools in the network may give you the views you want, but without access, the information is not readily available to you**

**Standalone CPU based tools are not expandable to view the outside components**

**Others will need to be involved as you delve into problems, but the tools at your disposal need to give you basic information in order to proceed**



## Key Elements in User Data

**Can you determine information on not only the CPU involved, but also other network components that may be impacting the problem?**

**Can you determine availability and response times for the users of your CPU?**

**Real time and historical data is needed..  
One to solve immediate problems... One to allow capacity planning**

**Are commands access provided as well as alerting to operational consoles**

**Can a new employee quickly learn the system?**



# What is the User Activity?

Users - jlee22:14000 - SYSADMIN

File Edit View Help

Physical

Page: 2 of 2

ENTERPRISE

- UNIX Systems
- Linux Systems
  - ellnx01
    - Linux
      - Capacity Usage Information
      - Disk Usage
      - Network
      - Process
      - System Information
      - Users**
      - Log Alerts
    - vmInx1
      - Linux

Physical

### Process User Information

Line	File System User name	Real Group name	Effective Group name	Saved Group name	File System Group name	Real User ID	Effective User ID	Saved User ID	File System User ID
1	root	root	root	root	root	0	0	0	0
2	root	root	root	root	root	0	0	0	0
3	root	root	root	root	root	0	0	0	0
4	root	root	root	root	root	0	0	0	0
5	root	root	root	root	root	0	0	0	0
6	root	root	root	root	root	0	0	0	0
7	root	root	root	root	root	0	0	0	0
8	root	root	root	root	root	0	0	0	0
9	root	root	root	root	root	0	0	0	0
10	root	root	root	root	root	0	0	0	0

### User Login Information

User Name	User Login PID	Line	Login Time	Idle Time	Hostname(From)
rhone	2350	pts/0	10/11/02 18:51:31	00:02	YELLOWSTONE.candle.com

### Total User Logins

32767

0

1

Ready Server Available. Users - jlee22:14000 - SYSADMIN

## It all Leads to Service Level Management

Threshold	WAN	LAN
<b>CPU</b>	<b>75-90%</b>	<b>75-90%</b>
<b>Link</b>	<b>80-90%</b>	<b>40-90%</b>
<b>Memory</b>	<b>50%</b>	<b>50%</b>
<b>Output Queue</b>	<b>200</b>	<b>25</b>
<b>Buffer Misses</b>	<b>Any</b>	<b>Any</b>
<b>Broadcast Vol</b>	<b>10/Sec</b>	<b>300/Sec</b>
<b>FECN/BECN</b>	<b>10/Sec</b>	<b>N/A</b>



## Summary

**You never solve performance problems.....  
.....You just keep moving them**

**The basic performance issues remain  
the same.....But QoS adds a new view**

**Emerging applications need higher  
levels of performance**

**Performance data readily available  
.....but the interpretation and action  
plans are lax**

**Complexity**

**Expect change and new ideas to emerge**

**Policy systems required to ease administration complexity**



धन्यवाद

Hind Hindi

多謝

Traditional Chinese

ขอขอบคุณ

Thai

Спасибо

Russian

Gracias

Spanish

Thank You

English

شكراً

Arabic

Obrigado

Brazilian Portuguese

Grazie

Italian

多谢

Simplified Chinese

Danke

German

Merci

French

நன்றி

Tami Tamil

ありがとうございました

Japanese

감사합니다

Korean



# Resources

- *Yahoo Group Teamrooms*      [groups.yahoo.com](http://groups.yahoo.com)      (*NetView, TBSM\_Users*)
- Tivoli software homepage -- <http://www-3.ibm.com/software/tivoli/>
- Tivoli UserGroups --[http://www-3.ibm.com/software/sysmgmt/products/support/Tivoli\\_User\\_Groups.html](http://www-3.ibm.com/software/sysmgmt/products/support/Tivoli_User_Groups.html)
- Tivoli Customer Portal --<https://www6.software.ibm.com/reg/tivoli/custport-l>
- Tivoli Education -- <http://www-3.ibm.com/software/tivoli/education>
- Tivoli Software Events -- <http://www-3.ibm.com/software/tivoli/news/events/>
- Tivoli Best Practices -- <http://www-3.ibm.com/software/tivoli/features/oct2002/best.html>
- IBM Link - <http://www.ibm.link.ibm.com/>.
- IBM Manuals - <http://w3.ehone.ibm.com/public/applications/publications/cgi-bin/pbi.cgi>.
- IBM Software for zSeries On Demand Events -<http://www-3.ibm.com/software/is/mp/s390/ondemand/>

